

INVERSE MULTIPLEXER

Field of the Invention

This invention relates to the field of packet transmission, and in particular to an inverse multiplexer for sending a stream of cells or packets onto a plurality of transmission links.

- 5 The invention is particularly applicable to asynchronous transfer mode (ATM) transmission.

Background of the Invention

- In ATM technology, a stream of cells must be transmitted over a physical link to a receiver, where the cell stream is demultiplexed into a series of virtual channels. If only
10 one link is used, for example, a DS1 link, the transmission rate is obviously limited by the physical bandwidth of the link. The ATM Forum has defined a standard, known as IMUX, inverse multiplexing, whereby a cell stream can be divided among several links so that the total bandwidth to the cell stream is the aggregate bandwidth of the links. In accordance with a defined protocol, the cells are transmitted in a round robin fashion over
15 the several links and assembled into a single cell stream at the far end.

The device for performing this operation is an Imux transceiver. It may conveniently be incorporated on a single chip. The Imux transmitter takes a stream of cells, and transmits them over a series of physical links, typically E1 or T1 access lines, and to a receiver which generates a corresponding stream of cells at the far end.

- 20 Typically an Imux will have sixteen output ports capable of accommodating a corresponding number of links. Increasing the number of ports results in an increase in the complexity of the chip, which is not justified in all situations. It is desirable to have a standardized chip rather than customize the chip for different markets.

- For example, a DS3 link includes 28 T1 links multiplexed together. In order to send cells
25 over a DS3 link, they must first be divided among the 28 T1 channels, but it does not necessarily follow that the IMA group will be chosen from the first sixteen links. A conventional sixteen port Imux cannot accommodate the remaining twelve links, yet a device with 28 ports has greatly increased complexity and also has a more limited market.

An object of the invention is to alleviate these problems.

Summary of the Invention

According to the present invention there is provided an inverse multiplexer device comprising an input port for receiving a stream of data packets, a plurality of output ports for connection to outgoing physical links, transmit buffers for preparing outgoing packets, an expansion port capable of receiving packets from said buffers and transferring them through a corresponding expansion port on another like inverse multiplexer to designated output links on the other inverse multiplexer, and a controller for outputting the data packets on a group of any of said links in accordance with an inverse multiplex protocol, whereby said multiplexers can be cascaded to increase the number of output links that can be accommodated.

The packets are typically ATM cells, although the invention would be applicable to other network protocols. Another application of the invention is the use of IMA over SHDSL and HDSL2 lines. These lines are gradually replacing T1 and E1 lines.

Typically, a transmitter and complementary receiver are integrated on the same chip to provide a bi-directional device.

The expansion ports are preferably connected in the form of a ring carrying time division multiplexed data between two or more IMA devices connected together in a daisy chain arrangement. The ring has a plurality of parallel lines, as well as some control lines, since desirably the expansion port connects to the output buffers where the data flows are parallel.

A switch is preferably provided to disconnect the output port normally associated with an transmit buffer when packets output from that buffer are to be sent out on an IMA link connected to the other inverse multiplexer device.

It will thus be seen that it is possible to configure an IMA group using the expansion port from another similar device.

An important advantage of this device arises when it terminates one or more DS3 links. The DS1 links that are demultiplexed from the DS3 link can be physically connected to different IMA-devices. Through the use of the expansion port, any DS1 (from any IMA device, up to 6 devices) can be a member of an IMA group on any of the IMA-16 devices. This is done without any limitation on the current IMA implementation. The full IMA

specification is still supported with CTC and ITC clocking mode. Without the expansion port, an external TDM switch terminating 28 or more DS1 links would be required with a limitation on the timing mode. Only the CTC mode could be used. Furthermore, no grooming at the DS1 level is required as the device permits connection of any DS1 link to any IMA group and no external circuitry is required .

The limit to the number of devices connected in the daisy chained ring is normally six, and the number of outside connections that can be addressed is limited to a total of 32. A DS1 link that is connected to an on-chip IMA group does not consume any of the 32 addresses on the ring, as it does not have to pass the ring as it is passed directly from the transmit buffer to the output port.

Brief Summary of the Drawings

The invention will now be described in more detail, by way of example only, with reference to the accompanying drawings, in which:-

Figure 1 is a functional block diagram of a transmitter in IMA mode;

Figure 2 is a block diagram of a receiver in IMA mode;

Figure 3 shows the ring address register;

Figure 4 shows a series of devices connected together through their expansion ports; and

Figure 5 is a block diagram of a ring controller..

Description of the Preferred Embodiments

Referring now to Figure 1, the IMA transmitter, which is also capable of operating in the UNI mode, i.e. with one cell stream being assigned to a single link, comprises a Utopia level 2 interface 10 for receiving an ATM cell stream from an ATM layer device, such as an ATM switch (not shown). The transmitter is normally integrated on single chip with the receiver shown in Figure 2.

An IMA (Inverse Multiplexer for ATM) Data Cell Rate Generator 12 controls the flow of cells through cell control block 14 into FIFO transmit buffers 16 connected through ICP cell modifiers and cell scrambling blocks 18, and parallel connections 20 to parallel serial converters 22 providing output ports for connection to respective links 24. The FIFO transmit buffers 16 serve as the TX UTOPIA FIFO when the link is configured in

the UNI mode (i.e. without the inverse multiplexing feature enabled) and the TX link FIFO when the device is configured in the IMA mode.

5 The device is under overall control of a microprocessor through interface 26. Controller 28 provides the round robin scheduler and FIFO selection and adaptive shaper to inverse multiplex the outgoing cell through blocks 18 and connections 20 onto outgoing links 24 after parallel to serial conversion in blocks 22.

In the UNI mode, i.e. without inverse multiplexing active, incoming cells are passed to link FIFOs 16 (link 0 to link 15) and out through the associated units 18 connected directly to the respective output ports 22.

10 In the IMA mode, the incoming cells are first passed to the group FIFOs 16 (group 0 to group 7), from where there are placed in the link FIFOs (link 0 to link 15) by the round robin scheduler 28 for transmission over the links forming the IMA groups. The device described so far operates in a conventional manner in accordance with the ATM forum specifications except that the link FIFOs can now be associated with physical links
15 present on a different device.

In accordance with the principles of the invention, the outputs of blocks 18 can be diverted through parallel connections 29 and ring 30 to links associated with a separate but similar device cascaded with the present device. The ring 30 consists of eight lines carrying parallel data as well as some additional control lines carrying control signals
20 between different devices. It thus acts as an interface to interconnect two or more similar devices and allow links connected to another device to form part of an IMA group with physical links on the first device. Each device has a ring control expansion port 33 connected in a daisy chain arrangement.

Switches 31 break the connection between the ports 22 and the blocks 18 when the cells
25 are to be diverted to a link attached to another device through the ring 30. Thus, for example, a cell leaving the first Unit 18a would be diverted from link 0 to the selected link in the cascaded device connected to the ring 30 through expansion port 33.

The receive circuitry is shown in Figure 2. Incoming DS1 links are connected to serial-to-parallel converters 40, cell delineation blocks 42, ICP processing blocks 44, IMA frame
30 state machines 46 and RAM controller 48. The ICP cell processing blocks are connected

to link info registers 50 and a buffer 52 storing ICP cells with changes. These are connected through microprocessor interface 54 to the controlling microprocessor (not shown).

5 The RAM controller 48 is connected to RAM 56, rate recovery block 58, RX scheduler 60 and UTOPIA interface 62. This interface outputs are stream of ATM cells corresponding to the input stream.

The operation of the receive circuitry described so far is conventional for an IMA receiver. In a similar manner to the transmitter, the links 62 between the serial-to-parallel converters 40 and their corresponding cell delineation blocks 42 can be broken by
10 switches 64, which permit incoming cells to be routed into a similar cascaded device through the common ring 30. For example, the top link $DST_{[0]}$ could form part of an IMA group with links from a similar cascaded device (not shown). Alternatively, cells arriving on a link, say $DST_{[3]}$ of a similar device could be routed through the ring 30 to the cell delineation block 42 for link $DST_{[0]}$ in the place of cells normally arriving on this link.

15 It will thus be seen how the device can be cascaded with other like devices through the ring to allow IMA groups to be formed of links connected to different devices. This increases the total number of links N from which a particular IMA group consisting of n links can be formed, where $n \leq N$, without the need for producing a custom device or significantly increasing the pin count. For example, any selection of 28 DS1 links
20 forming a DS3 trunk can be made to form an IMA group by cascading two devices together through the ring as described without the need for any external circuitry or switch.

Figure 4 shows three IMA devices 100 connected together in ring 30 by means of their respective expansion ports 33. Each expansion port has a transmit and receive side for
25 respectively transmitting data on and receiving data from the ring 30.

Figure 5 shows the expansion port 33 in more detail. A ring address register 70, shown in more detail in Figure 3, associated with each ICP cell modifier block 18 stores the address of the outgoing physical link. If this is the link 24 connected directly to the ICP modifier block 18, enable bits are set and cause direct connection controller 82 to close switch 31
30 in order to pass the output byte directly to the associated physical link. If the address

corresponds to a link on another device, switch 31 is opened and the Tx byte is assembled into a message in message assembler 84 for transmission over the ring 30. The message is first passed into Tx FIFO 86 and then ring driver 88, which physically places the message on the ring 30.

- 5 Byte Rx input to message assembler 84 comes from a serial-to-parallel to converter 40 on the receive side of the device. If the received byte is intended directly connected cell delineation block 42, direct connection controller closes switch 64. If the incoming link is associated with an IMA group on another chip, the received byte Rx is passed through message assembler 84, which adds the address of the destination link, and passed through
10 transmit FIFO 86 and ring driver 88 for transmission to the appropriate device over ring 30.

- Byte coming off the ring 30 are received in ring receiver 90 and passed to address comparator 92, wherein they are compared with addresses stored in ring address register 94. If an incoming message contains a byte intended for the device, it is output
15 form the comparator 92 either as a transmit byte Tx or a receive byte Rx. If it is a transmit byte Tx, it is passed to the appropriate output port 22. If it is a receive byte Rx, it is passed to the appropriate cell delineation block 42 of the receive portion of the device.

- The expansion port 33 is under overall control of ring controller 94, which automatically forwards messages not intended for the device to the transmit FIFO 86 for output to the
20 ring 30. One IMA device is designated the master, and the ring master control 94 for this device has the additional function of discarding redundant messages that have travelled around the ring. Before forwarding the messages, the device of the master has the additional function of performing a check to see whether they should be discarded.

- When two or more devices are cascaded together, they are first connected in a daisy chain
25 arrangement through their expansion ports 33. One device must be configured as the ring master. This is done by writing to a control register. Each TX Cell Modifier Block 18 and RX Cell Delineation Block 50 available to form part of an IMA group is assigned a ring address ranging from 0 to 31. The same ring address is used for both the TX Cell Modifier Block 18 and the associate RX Cell Delineation Block 50, but the same address
30 should not be used for more than one TX Cell Modifier 18/RX Cell Delineation block 50 pair.

Likewise, each TX port 22 and RX port 40 is also assigned a common ring address, but the same address should not be used for more than one TX/RX port 22, 40 pair.

When the outgoing packet is to be routed through the ring 30, then the paired TX and RX ports are assigned the same ring address. The connection is made by matching the ring address of the TX Cell Modifier Block 18 to the TX port 22, and matching the address of the RX port 40 to the RX Cell Delineation block 50.

There are 16 registers to assign a ring address to the or Cell Modifier/Cell delineation pair (18/42) and some additional bits in the control register.

The ring 30 can run at the system clock speed or at half of the frequency of the system clock. In the latter case, only 16 ports can be routed through the expansion port.

Although the TX and RX ports 22, 40 and the TX Cell Modifier 18 and RX Cell Delineation blocks 42 are paired, it is still possible to support the asymmetrical mode for an IMA group (where the number of TX and RX links can be different). In this case, the software does not enable either the TX port 22, or RX port 40.

It is possible to employ an enhanced switching mode where more than one fractional T1/E1 channel is merged internally into a selected port. This permits the multiplexing of multiple channels onto a single TDM link, which eliminates the use of an external switch.

The ring address register 70, shown in Figure 3, is a 16 bit register. Bits 12 to 8 of the register define the ring address of the TX Cell Modifier Block 18/RX block 50. Bits 4 to 0 define the TX/RX ring address. The bit 15 and 7 are enable bits which define if the ring connection is enabled for the TX/RX port or if the default internal direct connection is to be used. When bit 15 and bit 7 are 0, the normal connection between the cell modifier block and the Tx port 22, and between the Rx port 40 and cell delineation block 42 is established.

A ring connection can be made to another port that is on the same device but not directly connected to the cell modifier block 18. In this case the data is carried over the ring 30 to get to the destination port is on the same device.

The data packets are carried on the ring 30 in the form of two-byte messages as shown in Table 1 below. The first byte contains the address in the lower 5 bits (bit 0 to 4). The upper 2 bits are used to define the type of message, as defined in the table below.

5 The bit 5 is used as a maintenance bit, which is controlled ring master control 94 of the device which is defined to be the Ring Master. For each message entering the Master Ring device, if this bit is set to 1, the device replaces the message type by an empty message. Counting the number of occurrences Status Bits indicates a problem. When the bit is 0 in the message entering the Ring Master device, it is set to 1 and the rest of the message is not altered. If it goes through the ring without being addressed, it will get
10 discarded when coming back to the Ring Master .

The second byte is a data byte forming part of a data packet, typically an ATM cell, to be transmitted over a link. In this case 53 messages are need to transfer a cell to a link connected to different device.

Table 1: Message type

Bit 7: 6	Message Type
00	Empty Message
01	Request a byte
10	Byte Received
11	Byte to Transmit

15

In the UNI mode, that is without employing the inverse multiplexing option, a DS1 link, for example, can be terminated in any of the cascaded devices without passing through the ring 30.

20 In another application, a single 16-port IMA device would be able to support $n \times 64$ timeslots for use with SHDSL and HDSL2 lines. Typically at a frame rate of 8KHz, the number of timeslots can be programmed from 1 to 32 in steps of one timeslot, and then in multiples of 32 timeslots from 32 to 28 timeslots per frame.

25 It will thus be appreciated that the described device permits a number of physical links greater than the number of ports on any one physical device to be formed into an IMA group by cascading like devices together without the need for complex external circuitry. A the same time, the internal structure of the device is standard.